Yuwen (Leah) Jia

Mentor: Nina Galanter

SPA DRP Winter 2022

Optimal Treatment Rules

Causal Inference and Statistical Learning

Throughout this quarter, we explored the application of statistical learning in estimating optimal treatment rules to make decisions about whether or not to apply a treatment to a patient. Creating such a treatment rule is practical in situations where different individuals react to a particular treatment, for instance, a specific type of medication, differently.

The first few weeks were dedicated to an introduction to causal inference, including understanding the difference between causation and association and concepts of counterfactual outcomes and expectations. Three conditions must be satisfied in order to analyze causal inference—consistency, exchangeability, and positivity. Here we are dealing with situations where we want our treatment and outcomes to be well-defined. We also want to have equally distributed covariates (not confounding factors) for each treatment level and a positive probability of being assigned to each treatment level, both of which can be guaranteed by using randomization.

With the fundamental definitions in mind, we then extended these concepts to estimating an optimal treatment rule—a function that takes in the covariates and outputs a decision—and our goal is to maximize the average outcomes. We used the Q learning method. In this case, a particular linear mean model that explains the given dataset is needed, with appropriate parameters. Assuming we have a continuous outcome variable (***Y***) and a binary treatment variable (***A***), we are ready to approach optimal rules using a proposed model for the outcome ***Y***. We separate the model into two components—terms that relate solely to patient information ***X*** and terms that interact with treatment and thus inform our treatment decisions,

$$Y = \beta_0 + \beta_1 X + (\beta_2 + \beta_3 X)A \ .$$

This model offers a way to estimate the result for future observations or interpret the relationship between the output and the predictors. Here, we are only concerned with the terms related to the treatment, so we want to compute the last two parameters ($\boldsymbol{\beta}_2$, $\boldsymbol{\beta}_3$) in this model using standard linear regression. Once the parameters are obtained, we can then estimate an optimal treatment decision rule, that is to apply treatment ($\boldsymbol{A}$=1) if the treatment related term is greater than 0, and not apply the treatment ($\boldsymbol{A}$=0) otherwise. This Q-learning method is one of many approaches to estimate optimal treatment rules.

The last part of the project was to test and evaluate the methods we have learned, where we used simulations with data generated from random sampling. We generated one dataset for fitting and training the model, and another for testing the performance of the methods. Based on a linear model like the one shown above, we can build a treatment decision rule. We compared the performances of three different rules based on three separate models: a correctly specified model, a misspecified model, and the true model. By comparing the accuracy of our predicted decisions from the three rules, we can see how the methods perform and which of them are better than the others at predicting optimal treatment decisions. In general, simulating data gives us the opportunity to visualize the performances of our methods.

In conclusion, the q-learning method helps us figure out the optimal treatments that maximize average outcome values across all individuals, though the significant differences in average accuracy between using a correctly specified model and an incorrectly specified model proves that we should be cautious when proposing a model if we want a certain level of accuracy. Nevertheless, in most cases where the truth is unknown, the q-learning method is a particularly useful tool for determining optimal treatments.

Sources and references:

Chapter 1-3, 11, 15, *Causal Inference: What If* by Miguel A. Hernán and James M. Robins

Chapter 2, *An Introduction to Statistical Learning* by Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani

Chapter 6: Analysis in the single-stage setting: An overview of estimation approaches for dynamic treatment regimes, *Adaptive Treatment Strategies in Practice*, Michael P. Wallace and Erica E. M. Moodie